Automatic Feature Extraction for Panchromatic Mars Global Surveyor Mars Orbiter Camera Imagery

Catherine S. Plesko [a*], Steven P. Brumby [b], and Conway Leovy [c]

[a] Department of Physics, University of Washington
[b] Space and Remote Sensing Sciences, Los Alamos National Laboratory,
Mail Stop D436, Los Alamos, New Mexico 87545
[c] Department of Atmospheric Sciences, University of Washington,
Box 351640, Seattle, Washington 98195

The Mars Global Surveyor Mars Orbiter Camera (MOC) has produced tens of thousands of images, which contain a wealth of information about the surface of the planet Mars. Current manual analysis techniques are inadequate for the comprehensive analysis of such a large dataset, while development of handwritten feature extraction algorithms is laborious and expensive. This project investigates application of an automatic feature extraction approach to analysis of the MOC narrow angle panchromatic dataset, using an evolutionary computation software package called GENIE.

GENIE uses a genetic algorithm to assemble feature extraction tools from low-level image operators. Each generated tool is evaluated against training data provided by the user. The best tools in each generation are allowed to "reproduce" to produce the next generation, and the population of tools is permitted to evolve until it converges to a solution or reaches a level of performance specified by the user. Craters are one of the most scientifically interesting and most numerous features in the MOC data set, and present a wide range of shapes at many spatial scales. We now describe preliminary results on development of a crater finder algorithm using the GENIE software.

Keywords: Feature Extraction, Genetic programming, Supervised classification, Mars, Craters, Panchromatic imagery.

1. INTRODUCTION

NASA's Mars Global Surveyor[1] (MGS) orbiter spacecraft has been studying Mars since 1997. It carries a variety of scientific instruments including the Mars Orbiter Camera[2] (MOC). The MOC is made up of two wide-angle cameras and one narrow-angle, high-resolution camera. The narrow angle dataset, which is the focus of this study, provides imagery with a spatial resolution generally of the order of 3 meters/pixel. Since arriving at the planet, the MOC has taken over 80,000 images which have been used to study geologic and atmospheric processes, including sediments, dust streaks, and volcanism (e.g., see Refs. 1, 3-6). MGS completed its primary mission on January 31, 2001, and the spacecraft continues to take data on areas of interest.

There have been several databases developed to aid in the processing of Mars data. The Mosaicked Digital Image Model (MIDIM) and the USGS three-dimensional Digital Terrain Model (DTM) combined Viking orbiter data sets were the first digital atlases of Mars. The Planetary Data System[10] (PDS) and Malin Space Sciences web site[11] maintain on-line databases of MOC data which are sorted by orbit number and image number within each orbit. They also provide a map-based browsing interface.

There are a variety of terrain features in MOC images that are interesting to researchers. For example, craters are common and a highly useful surface feature in MOC data. They can be used to determine relative ages of the surface and give clues about its composition (see, e.g., Refs. 7-9). Sand dunes show the flow patterns of the winds that formed them, give clues to dust particle size and abundance, and are also useful in terrain age estimates. Landslides, dust devil tracks, and "yardangs" (a large-scale wind-erosion feature usually found in exposed rock), give evidence of various types of surface weathering processes. It would be a great aid to research if a database was available such that researchers could set search parameters based on the surface features in the scene. However, the sheer volume of data makes traditional photo-interpretation and manual analysis of the MOC images for such a database extremely time consuming.

[*] Correspondence: Email: csp@u.washington.edu

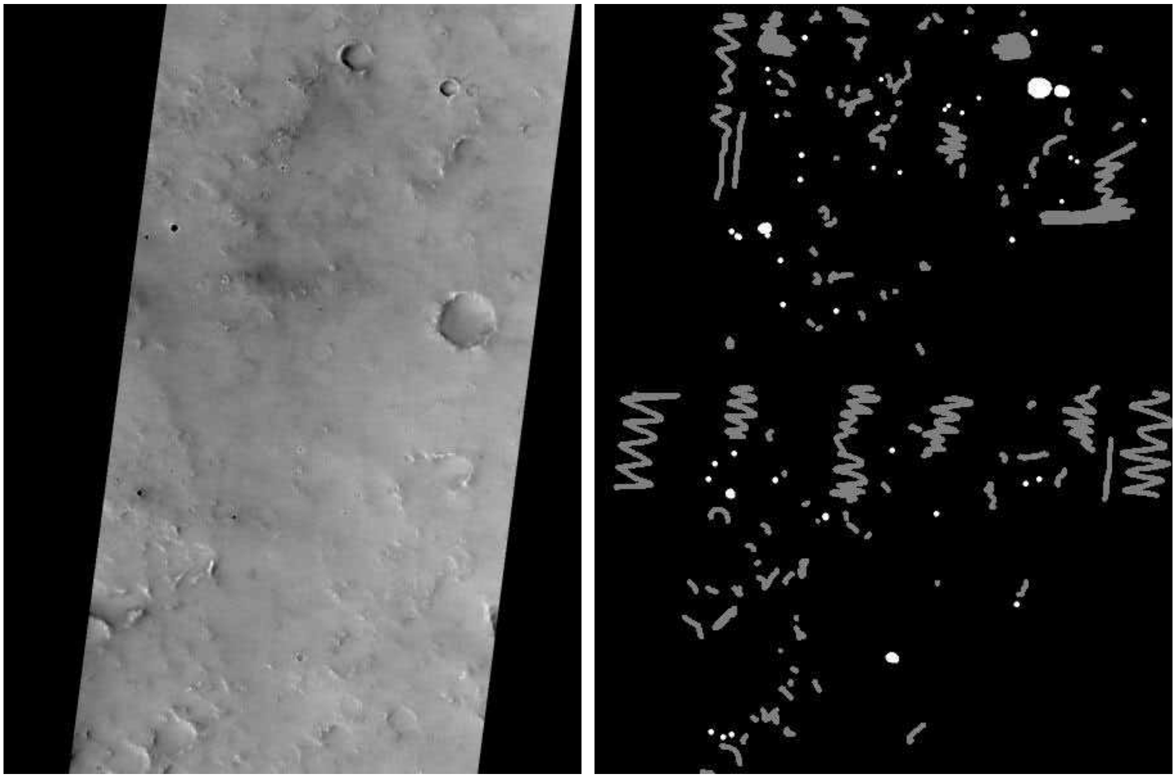

Figure 1. The training scene (Left) is the first thousand pixel rows of the image M1500956, an intercrater plane in the Deucalionis Regio area of the Martian southern hemisphere. The system was trained with a truth file (Right) that marked selected fresh bowl shaped craters as true (shown in white) , and selected ridges and non-cratered terrain as false (gray).

In order to automate the creation of a feature database, a set of feature-extraction algorithms is needed. These algorithms could then be applied to the dataset in batches. The results could then provide the content for a feature database. Given the need to develop a number of such algorithms reasonably quickly, machine learning provides an interesting alternative to human-design of algorithms. Such an approach has the potential to shorten the time required to produce a range of automatic feature extraction tools, and could enable the construction of the feature database envisioned.

For the present study we will describe the development of a crater finding algorithm, using a genetic programming approach originally developed for Earth remote sensing applications. We are initially interested in developing a pixel-level algorithm that will segment out spatial regions likely to contain craters, and in future work will explore higher-level, scale-invariant processing which uses the geometry of the segmented regions (e.g., linear feature, circular feature, annular feature) to complete the classification of craters as geometric objects.

2. AUTOMATIC FEATURE EXTRACTION SOFTWARE: GENIE

GENIE[12,13,14,15] is an evolutionary computation (EC) software system which uses a genetic algorithm[16,17,18] (GA) to assemble image-processing algorithms from a collection of low-level image processing operators (e.g., edge detectors, texture measures, spectral operations, and various morphological filters). This system has been shown to be effective in looking for complex terrain features, such as wildfire burn scars[19]. GENIE has been described at length elsewhere[12,20], so we will only present a brief description of the system here.

GENIE follows the classic evolutionary paradigm: a population of candidate image-processing algorithms is randomly generated, and the fitness of each individual assessed from its classification performance on a user provided training scene.

After fitness has been assigned, the most fit members of the population are permitted to reproduce with modification via the evolutionary operators of selection, crossover, and mutation. The process of fitness evaluation and reproduction with modification is iterated until some stopping condition is satisfied. We now briefly describe our method of providing training data, our encoding of image-processing algorithms as chromosomes for manipulation by the GA, and our method for evaluating the fitness of individuals in the population.

2.1 Training Data

The environment for the population consists of one or a number of training scenes. Each training scene contains a raw image, together with a weight plane and a truth plane. The weight plane identifies the pixels to be used in training, and the truth plane locates the features of interest in the training data. Providing sufficient quantities of good training data is crucial to the success of any machine learning technique. In principle, the weight and truth planes may be derived from an actual ground campaign (i.e., collected on the ground at the time the image was taken), may be the result of applying some existing algorithm, and/or may be marked-up by hand using the best judgement of an analyst looking at the data. We have developed a graphical user interface (GUI), called Aladdin, for manual marking up of raw imagery. Using Aladdin, the analyst can view a raw image in a variety of ways, and can mark up training data by painting directly on the image using the mouse. Training data is ternary-valued, with the possible values being “true”, “false”, and “unknown”. True defines areas where the analyst is confident that the feature of interest does exist. False defines areas where the analyst is confident that the feature of interest does not exist. Unknown pixels do not influence the fitness of a candidate algorithm.

2.2 Representation of Image-Processing Algorithms

Traditional genetic programming[21] (GP) uses a variable sized (within limits) tree representation for algorithms. Our representation differs in that it allows for reuse of values computed by sub-trees, i.e. the resulting algorithm is a graph rather than a tree. The image processing algorithm that a given chromosome represents can be thought of as a directed acyclic graph where the non-terminal nodes are primitive image processing operations, and the terminal nodes are individual image planes extracted from the multi-spectral image used as input. Our representation also differs in that the total number of nodes is fixed (although not all of these may actually be used in the final graph), and crossover is carried out directly on the linear representation.

We have restricted our “gene pool” to a set of useful primitive image processing operators (“genes”). These include spectral, spatial, logical and thresholding operators. The set of morphological operators is restricted to function-set processing morphological operators, i.e., gray-scale morphological operators having a flat structuring element. The sizes and shapes of the structuring elements used by these operators is also restricted to a pre-defined set of primitive shapes, which includes the square, circle, diamond, horizontal cross and diagonal cross, and horizontal, diagonal, and vertical lines. The shape and size of the structuring element are defined by operator parameters. Other local neighborhood/windowing operators such as mean, median, etc., specify their kernels/windows in a similar way. We define scratch planes as blocks of memory for storing intermediate calculations within a candidate image-processing algorithm. Once “scratch” planes have been generated, GENIE is allowed to explore weighted sums, differences and ratios of data and scratch planes.

A single gene consists of an operator, plus a variable number of input arguments specifying from where input is read, output arguments specifying where output is to be written, and any additional parameters that might be required to specify how the specific operator works (e.g., the diameter and shape of a structuring element used in a morphological filter). The operators used in Genie take one or more distinct image planes as input, and generally produce a single image plane as output. Input can be taken from any data plane in the training data image cube. Output is written to one of a number of scratch planes, temporary workspaces where an image plane can be stored. Genes can also take input from scratch planes, but only if that scratch plane has been written to by another gene positioned earlier in the chromosome sequence. We use a notation for genes [10] that is most easily illustrated by an example: the gene `[ADDP rD0 rS1 wS2]` applies pixel-by-pixel addition to two input planes, read from data plane 0 and from scratch plane 1, and writes its output to scratch plane 2. Any additional required operator parameters are listed after the output arguments.

Note that although all chromosomes have the same fixed number of genes, the effective length of the resulting algorithm may be smaller than this. For instance, an operator may write to a scratch plane that is then overwritten by another gene before anything reads from it. GENIE performs an analysis of chromosome graphs when they are created and only carries out those processing steps that actually affect the final result. Therefore, the fixed length of the chromosome acts as a maximum effective length.

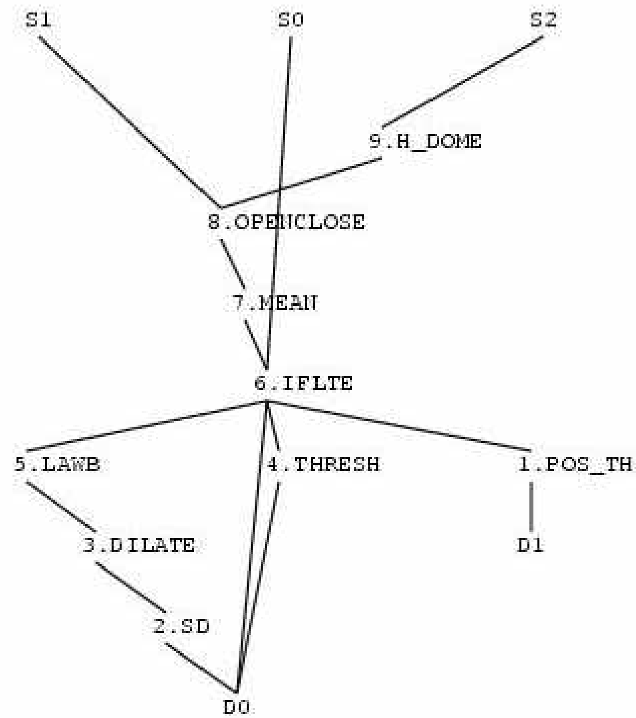

Figure 2. Graph of evolved crater finding algorithm: Each node (gene) is labeled by its position along the length of the chromosome and by the GENIE software's mnemonic for the primitive image processing operator (e.g., 1. POS_TH).

2.3 Supervised Classification and Fitness Evaluation

Each candidate image-processing algorithm generates a number of intermediate feature planes (or “signature” planes), which are then combined to generate a Boolean-valued mask for the feature of interest. This combination is achieved using a standard supervised classifier (we use the Fisher linear discriminant[22]), and an optimal threshold function.

Complete classification requires that the image-processing algorithm produce a binary-valued output plane for any given scene. It is possible to treat, e.g., the contents of the first scratch plane as the final output for that candidate image- processing algorithm (thresholding would generally be required to obtain a binary result, though Genie can choose to apply its own Boolean thresholding functions). However, we have found it to be useful to perform the combination of the data and scratch planes using a non-evolutionary method, and have implemented a supervised classifier backend. To do this, we first select a subset of the scratch planes and data planes to be “signature” planes. For the present experiments, this subset consists of just the scratch planes. We then use the provided training data and the contents of the signature planes to derive the Fisher Discriminant, which is the linear combination of the signature planes that maximizes the mean separation in spectral terms between those pixels marked up as “true ”and those pixels marked up as “false”, normalized by the total variance in the projection defined by the linear combination. The output of the discriminant-finding phase is a real-valued single-plane “answer” image. This is reduced to a binary image by exhaustive search over all the training pixels to find the threshold value that minimizes the total number of misclassifications (false positives plus false negatives) on the training data.

The fitness of a candidate solution is given by the degree of agreement between the final binary output plane and the training data. This degree of agreement is determined by the Hamming distance between the final binary output of the algorithm and the training data, with only pixels marked as true or false (as recorded in the weight plane) contributing towards the metric. The Hamming distance is then normalized so that a perfect score is 1000.

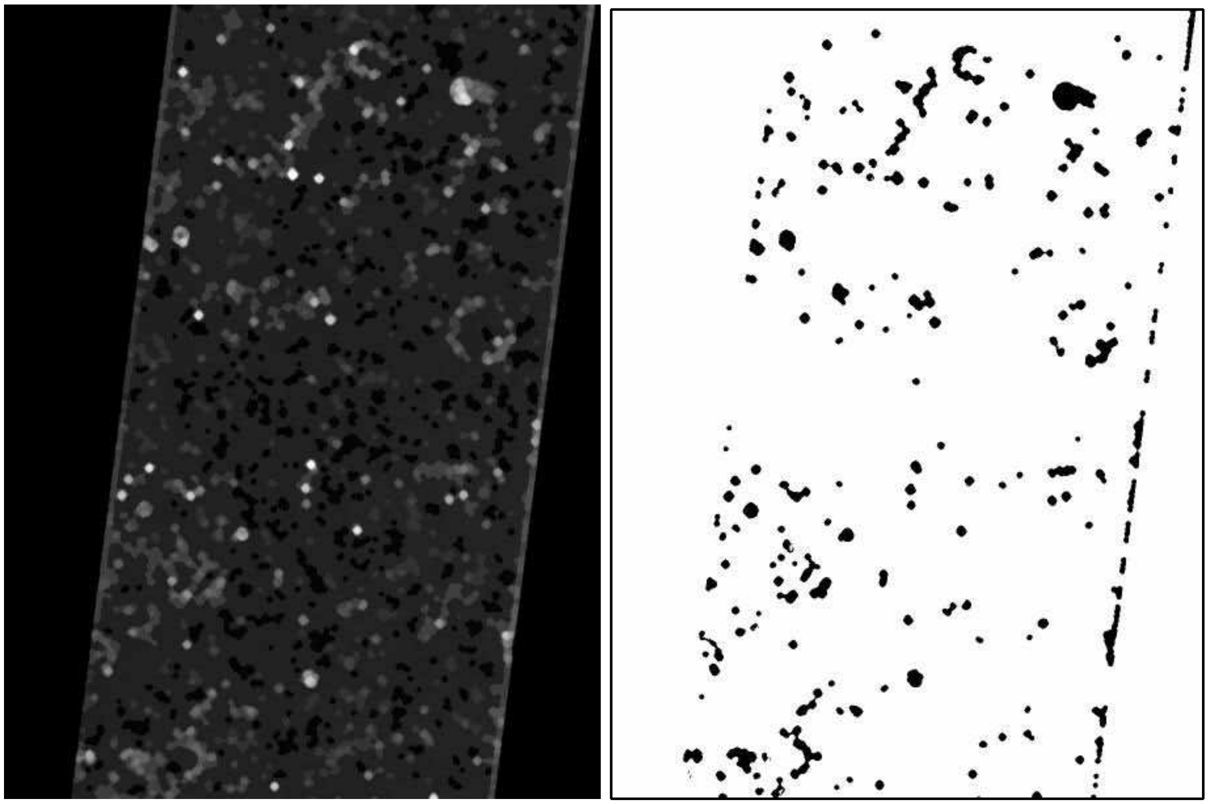

Figure 3. Grayscale answer plane (Left) and Boolean answer plane (Right, “crater” pixels marked in black) for the algorithm shown in Figure 2 and the training image shown in Figure 1.

3. RESULTS

For the present study, we chose to apply the GENIE software to the task of finding craters in Mars Orbital Camera narrow-field observations. We selected two images: one image (Fig. 1) to train GENIE on, and one image to test the resulting evolved image processing algorithms. The training image (MOC image M1500956, Ref. 23) was chosen to present a reasonably homogeneous terrain marked by a number of bowl shaped craters obvious to the human eye. GENIE was trained on the first thousand pixel rows of the image with a truth file based on manual analysis (Fig. 1) that marked some of the fresh, bowl shaped craters in the scene as true, and some of the protruding surface features and non-cratered terrain as false. GENIE was run with a population of 20 chromosomes per generation for 477 generations. Each candidate algorithm contained a maximum of twelve image processing operations, and used three intermediate scratch planes. The best crater finding algorithm achieved a score of 936, with a detection rate of 95% and a false alarm rate of 8%. The algorithm found is represented in our encoding (see section 2.2 and Refs. 12, 20) as:

[POS_TH rD1 wS1 2 1][SD rD0 wS0 2 1][DILATE rS0 wS0 2 1][THRESH rD0 wS2 0.14]
[LAWB rS0 wS0][IFLTE rS1 rD0 rS0 rS2 wS0][MEAN rS0 wS1 3 1]
[OPENCLOSE rS1 wS1 4 1][H_DOME rS1 wS2 53]

We now describe the operation of the key “genes” in this image processing algorithm, which is also shown as a graph in Figure 2.

1. [SD rD0 wS0 2 1] The Standard Deviation in a 5 pixel diameter circular neighborhood of each pixel is calculated, which provides a measure of local variance, and is stored in scratch plane S0.

2. [DILATE rS0 wS0 2 1] A Dilation of the scratch plane S0 with a flat circular structuring element of diameter 5 pixels.

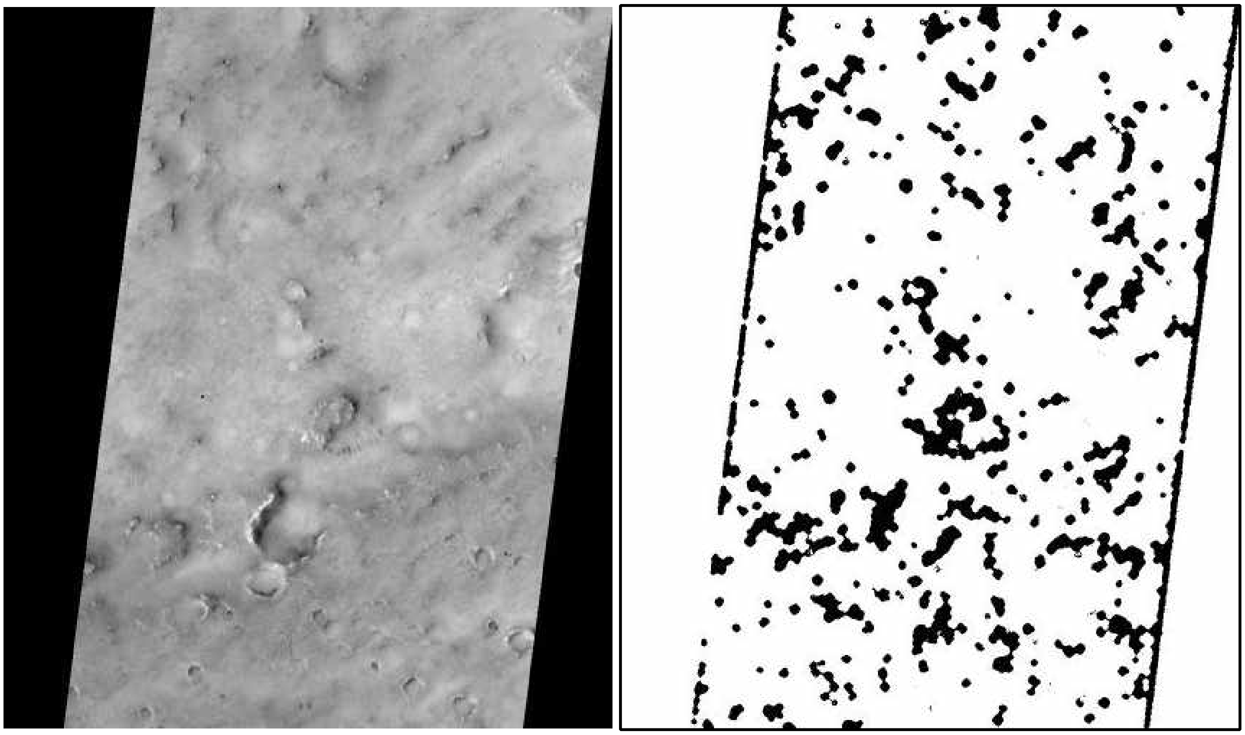

Figure 4. The algorithm was then tested on the next 1,000 pixel rows of M1500956 (Left), which it had not seen during its evolution. It was successful at finding the craters in the scene, but had difficulty distinguishing between concave and convex terrain.

3. [LAWB rS0 wS0] A local texture measure, defined by Law[24].

4. [MEAN rS2 wS1 3 1] Local mean in a 7 pixel diameter circular neighborhood.

5. [OPENCLOSE rS1 wS1 4 1] A morphological Open/Close operation, using a 9 pixel diameter, flat circular structuring element.

6. [H_DOME rS1 wS2 53] A standard morphological image processing operation, also known as "regional maxima"[25], where the parameter value of 53 is used to calculate the "floor" value for the regional maximum.

The final answer is formed from a linear combination of the data plane with the results of steps 5 and 6, and is thresholded to produce a simple Boolean classification for each pixel. The threshold value was chosen to minimize the classification error, that is, to minimize the total number of false positives and false negatives. The results of this algorithm on its training scene are shown in Fig. 3, before and after thresholding.

On inspection of this result, it appears that GENIE was quite successful at finding the craters within the training image, although it detected some protruding terrain as well (e.g. mesas, knobs). This is reasonable, as GENIE was not given any direct information regarding the direction of illumination. Although the ring of a crater and the walls of a mesa can look very similar, the shadows they cast will be on opposite sides of the object relative to the sun. It may be possible to correct for this with a filter that would take the angle of the sun and position of shadows relative to the object in interest.

The evolved algorithm is a stand-alone program that can be run on other images. We tested its out-of-training-sample behavior by applying it to the next thousand pixel rows of M1500956 (Figure 4), which it had not seen during its evolution. The result, shown in Fig. 4, is similar to its behavior throughout the training scene, which suggests a reasonably robust algorithm has been evolved.

4. CONCLUSIONS

The surface of Mars has been observed from orbit at a number of progressively finer spatial scales, from kilometers (Mariner series), to 100's of meters (Viking), and most recently at the few meter scale (MGS). It is this new data set that holds out the greatest chance of significant progress in understanding the history of the Martian surface, while also presenting the greatest challenge to traditional manual analysis techniques. This study investigated the evolution of a crater finding algorithm for application to the Mars Orbiter Camera narrow angle dataset. We described the algorithm, its application, and compared its results on training and test images to manual classification of the scenes. The algorithm is successful at detecting craters within the images, and generalized well to an image that it had not seen before. Although it currently detects convex features (e.g. mesas) as well as craters, this can in principle be improved with a filter that checks the position of shadows relative to the object and by structural analysis of the segmented regions. We find these results to be encouraging for the application of GENIE to the MOC panchromatic dataset.

ACKNOWLEDGEMENTS

GENIE research and development is funded by the U.S. Departments of Defense and Energy. S.P.B. would like to thank his colleagues on the GENIE software team: J.J. Bloch, M. Galassi, N.R. Harvey, S. Perkins, R.B. Porter, J.J. Szymanski, J. Theiler, and A.C. Young. C.S.P. acknowledges the support of a Mary Gates Foundation research great, and the financial support of the Washington NASA Space Grant Consortium. C.S.P. also wishes to thank J. Armstrong and T. Tavener for useful discussions.

REFERENCES

1. A. L. Albee, F. D. Palluconi, R. E. Arvidson, "Mars Global Surveyor Mission: Overview and Status", Science, Vol. 279, p. 13, 1998.
2. M. C. Malin, G. E. Danielson, A. P. Ingersoll, H. Masursky, J. Veverka, M. A. Ravine, T. A. Soulanille, "Mars Observer Camera", Journal of Geophysical Research, Vol. 97(E5), pp. 7699-7718, 1992.
3. K. S. Edgett, and M. C. Malin, "New views of Mars eolian activity, materials, and surface properties: Three vignettes from the Mars Global Surveyor Mars Orbiter Camera," Journal of Geophysical Research, Vol. 105(E1), pp.1623-1650, 2000.
4. W. K. Hartmann, M. Malin, A. McEwen, M. Carr, L. Soderblom, P. Thomas, E. Danielson, P. James, and J. Veverka, "Evidence for recent volcanism on Mars from crater counts," Nature, Vol. 397, pp. 586-589, 1999.
5. P. C. Thomas, M. C. Malin, M. H. Carr, G. E. Danielson, M. E. Davies, W. K. Hartmann, A. P. Ingersoll, P. B. James, A. S. McEwen, L. A. Soderblom, and J. Veverka, "Bright dunes on Mars," Nature, Vol. 397, pp. 592-594, 1999
6. J.R. Zimbelman, "Non-active dunes in the Acheron Fossae region of Mars between the Viking and Mars Global Surveyor eras", Geophysical Research Letters. Vol.27, no.7, pp.1069-1072, 2000.
7. M. H. Carr, M. C. Malin, "Meter-scale characteristics of Martian channels and valleys", Icarus., Vol.146, no.2, p.366-386, 2000.
8. Hartmann, W.K.,"Martian cratering. VI. Crater count isochrons and evidence for recent volcanism from Mars Global Surveyor", Meteoritics and Planetary Science. vol.34, no.2, p.167-77, 1999.
9. Hartmann, W.K., Berman, D.C., "Elysium Planitia lava flows: crater count chronology and geological implications", Journal of Geophysical Research. vol.105, no.E6, p.15011-25, 2000.
10. S.K. McMahon, "Overview of the Planetary Data System", Planetary and Space Science. vol.44, no.1, p.3-12, 1996.
11. Malin Space Science Systems Mars Orbiter Camera (MOC) image archive web site: `http://www.msss.com/mars_images`
12. S.P. Brumby, et al., "Investigation of feature extraction by a genetic algorithm ", Proc. SPIE, Vol. 3812, pp. 24-31,1999; see also the GENIE web site `http://www.daps.lanl.gov/genie` .
13. J. Theiler, et al., "Evolving retrieval algorithms with a genetic programming scheme ", Proc. SPIE, Vol. 3753, pp.416-425,1999.
14. N.R. Harvey, et al., "Parallel evolution of image processing tools for multispectral imagery", Proc. SPIE, Vol. 4132, pp. 72-82, 2000.
15. S. Perkins, et al., "GENIE: A hybrid genetic algorithm for feature classification in multi-spectral images", Proc. SPIE, Vol. 4120, pp 52-62, 2000.

16. J. H. Holland, Adaptation in Natural and Artificial Systems, University of Michigan, Ann Arbor (1975).
17. I. Rechenberg, Evolutionsstrategie: Optimierung technischer Systeme nach Prinzipien der biologischen Evolution, Fromman-Holzboog, Stuttgart (1973).
18. L. Fogel, A. Owens and M. Walsh, Artificial Intelligence through Simulated Evolution, Wiley, New York (1966).
19. S. P. Brumby, N. R. Harvey, J. J. Bloch, J. Theiler, S. Perkins, A. C. Young, and J. J. Szymanski, "Evolving forest fire burn severity classification algorithms for multi-spectral imagery". To appear in Proc. SPIE 4381. Presented at Aerosense 2001: SPIE's 15th Annual International Symposium on Aerospace/Defense Sensing and Controls, Orlando, Florida, April 16-20, 2001.
20. N. R. Harvey, J. Theiler, S. P. Brumby, S. Perkins, J. J. Szymanski, J. J. Bloch, R. B. Porter, M. Galassi, and A. C. Young,, "Image Feature Extraction: GENIE vs Conventional Supervised Classification Techniques", submitted to IEEE Transactions on Geoscience and Remote Sensing, 2001. Preprint: `http://www.daps.lanl.gov/genie`.
21. J. R. Koza, Genetic Programming: On the Programming of Computers by Natural Selection, MIT, Cambridge (1992).
22. For example, see C.M. Bishop, Neural Networks for Pattern Recognition, pp.105 –112, Oxford University (1995).
23. M.C. Malin, K. S. Edgett, M. H. Carr, G. E. Danielson, M. E. Davies, W. K. Hartmann, A. P. Ingersoll, P. B. James, H. Masursky, A. S. McEwen, L. A. Soderblom, P. Thomas, J. Veverka, M. A. Caplinger, M. A. Ravine, T. A. Soulanille, and J. L. Warren, “Intercrater plain in Deucalionis Regio”, NASA's Planetary Photojournal (http://photojournal.jpl.nasa.gov/), Catalog number M1500956, Created 5-14-2000.
24. K.I.Laws, “Rapid texture identification”, Proc. SPIE, Vol. 238, pp. 376-380, 1980.
25. L. Vincent, “Morphological grayscale reconstruction in image analysis: applications and efficient algorithms”, IEEE Trans. Image Processing, Vol. 2, pp. 176-201, 1992.